arXiv:1512.05610v2 [stat.ML] 7 Jun 2016

Classification of weak multi-view signals by sharing factors in a mixture of Bayesian group factor analyzers

Sami Remes[1], Tommi Mononen[1,2], and Samuel Kaski[1]

[1] Helsinki Institute for Information Technology HIIT, Department of Computer Science, Aalto University
[2] Department of Neuroscience and Biomedical Engineering, Aalto University

Abstract. We propose a novel classification model for weak signal data, building upon a recent model for Bayesian multi-view learning, Group Factor Analysis (GFA). Instead of assuming all data to come from a single GFA model, we allow latent clusters, each having a different GFA model and producing a different class distribution. We show that sharing information across the clusters, by sharing factors, increases the classification accuracy considerably; the shared factors essentially form a flexible noise model that explains away the part of data not related to classification. Motivation for the setting comes from single-trial functional brain imaging data, having a very low signal-to-noise ratio and a natural multi-view setting, with the different sensors, measurement modalities (EEG, MEG, fMRI) and possible auxiliary information as views. We demonstrate our model on a MEG dataset.

Keywords: Bayesian group factor analysis, brain decoding, MEG, multi-view learning, variational Bayesian inference

1 Introduction

Recently a lot of focus in machine learning has been given to the analysis of multi-view data, a scenario in which the data come from multiple data sources. One observation or data point consists of data from multiple sources, and in this sense each data source can be considered a different "view" to the same "object". This is a natural setting with, e.g. brain imaging data where each sensor, measurement modality (including EEG, MEG and fMRI) and possible auxiliary or experimental information form the views. Notably, the relevant, discriminating signal between, e.g., different experimental conditions is very weak compared to all other activity on-going in the brain.

There has been increasing interest to explore neuroscientific data using machine learning methods, which are more capable to reveal the complex phenomena happening in the brain than traditional contrasting methods (e.g. t-tests). A popular application area of machine learning is called brain decoding. The goal is to infer, based on a given brain signal, what task a subject is performing, given a training set consisting of sample signals of performed tasks. These samples are usually from single trials that have a low signal-to-noise ratio.

However, good prediction performance does not necessarily imply increase of neuroscientific knowledge about the processes in the brain. Some of the best prediction methods, such as SVM and Gaussian processes, are essentially black boxes, and it is hard to infer what the predictions are based on. Neuroscientists are interested in learning what parts of the brain have an effect on the prediction. Therefore, prediction methods that do not give these plausible explanations are much more difficult to use for accumulating neuroscientific knowledge. Generative probabilistic models are generally more immediately interpretable [11,7], and for instance when using linear generative models, the weights point to those brain locations that are activated due to a performed task.

A MEG device (as well as EEG) outputs a very fine-scale time series. Some methods such as the linear discriminant analysis and its various regularized versions that are widely used in brain-computer interfaces require that this rich source of information is squeezed into a single value (e.g. maximum or mean) [14,8]. Methods should take advantage of the time structure, otherwise useful information might get lost altogether or attributed to wrong time points.

We propose a novel Bayesian multi-view generative classification model that takes the time information into account in an intuitive way, by directly modelling the time dependencies. The generative approach is generally beneficial in scenarios where training data is scarce, because the assumptions of the generative process help learning the data efficiently. The model is motivated by setups common in brain imaging data, but applicability of the proposed model is not restricted only to neuroscience. The model can be applied to any time-varying or otherwise structured multiple-data-source framework.

Recently, unsupervised multiple-data-source modelling in the Bayesian framework has been studied, a state-of-the-art solution being the group factor analysis (GFA) model [18,13]. Similar models have been studied also with different priors for sparsity [22] and using an optimization-based sparse dictionary coding approach [9]. We propose learning a number of GFA models per class label, a mixture modelling approach commonly used for classification (e.g., in voice recognition with GMMs [15]). We jointly learn a set of clusters and their respective label distributions rather than learning models separately for each label. Additionally, we assume that the GFA models share some of their factors, which turns out very significant for classification. The large common parts of the signals can be modelled and explained away with the shared factors, while the weak and likely the most interesting discriminating parts of the signal can be modelled by the cluster-specific parameters.

Other approaches to multi-view learning also exist. A CCA-type approach was used where samples in multiple views are transformed to a common discriminative space where within-class variation is minimized while between-class variation is maximized for all views [10]. Approaches in [5] and [6] find a common latent space for the views using Gaussian processes. Optimization-based approaches [19,21] regress the data to the clusters and class labels, but they are not doing generative modelling in a way related to our GFA based approach. In [16] the authors studied using different types of features and applied classifiers with all the different subsets of the feature types, and showed that using a combination of features improved classification accuracy compared to a single feature type.

In the experimental section, we apply our model on MEG data and compare it against a group LASSO classifier [20] that is known to generally perform well, to the extent of being hard to beat in practice. The classifier takes the multi-source nature of the data into account in the same way as our model does, making the Group LASSO a natural baseline.

2 Model

First we present a GFA model for multi-view data, containing multiple data sources of possibly different dimensionalities. Next we extend GFA to a mixture model where some factors are shared between the mixture components (clusters). This effectively separates the parts of data that contribute to differences between the classes from the rest, allowing to share statistical strength to learn the shared parts more efficiently. Finally, the model is extended to include class labels.

2.1 Group Factor Analysis

Group factor analysis [13,18] can be seen both as an extension of Bayesian CCA [12] to multiple data sources, or alternatively it can be seen as an extension of factor analysis that treats the data sources similarly to how regular factor analysis treats individual variables.

Let $\mathbf{x}_n^{(m)} \in \mathbb{R}^{D_m}$ denote the nth sample (out of the data size N) of the mth data source, the model likelihood is given by

$$\mathbf{x}_n^{(m)} | \mathbf{W}, \mathbf{z}_n, \tau_m \sim \mathcal{N}(\mathbf{W}^{(m)} \mathbf{z}_n, \tau_m^{-1} \mathbf{I}) \; . \tag{1}$$

The observations (data points) are modeled with unknown latent variables $\mathbf{z}_n \in \mathbb{R}^K$ corresponding to K factors, which are then mapped to the data space with data source specific linear projections $\mathbf{W}^{(m)}$. The latent variables are shared between all data sources and can therefore model correlations between them, via controlling the values of $\mathbf{W}^{(m)}$ with a specific type of sparsity constraint; if some factor k is not useful for modelling the source m, we want $\mathbf{w}_k^{(m)}$ (the kth column of $\mathbf{W}^{(m)}$) to be zero. The desired structure is achieved with the automatic relevance determination (ARD) prior

$$\mathbf{w}_k^{(m)} | \alpha_k^{(m)} \sim \mathcal{N}(\mathbf{0}, (\alpha_k^{(m)})^{-1} \mathbf{I}), \tag{2}$$

where the α-parameters have independently non-informative gamma priors $\alpha_k^{(m)} \sim \Gamma(10^{-14}, 10^{-14})$.

2.2 Mixture of GFAs

We next describe a mixture of GFAs model that will be further extended later in this section. In our application scenario we assume that a single GFA model could explain most of the variation in the data, but not the most interesting part which is assumed to be weak. We assume the weak signal to consist of distinct parts (clusters), each of which can also be modeled by a set of factors; the clusters can consist of experimental conditions, for instance. The data is therefore generated by both cluster-specific factors $\mathbf{W}_c$, as well as shared factors $\hat{\mathbf{W}}$ (active regardless of cluster assignment) as follows:

$$\mathbf{x}_n^{(m)} | c_n \sim \mathcal{N}(\mathbf{W}_{c_n}^{(m)} \mathbf{z}_n + \hat{\mathbf{W}}^{(m)} \hat{\mathbf{z}}_n, (\tau_{c_n}^{(m)} \mathbf{I})^{-1}), \tag{3}$$

where a latent variable c_n indexes the cluster which the data point belongs to. The first term in the normal mean specifies the cluster-specific parts of data and the latter term the parts shared by the clusters (information not discriminating classes). Factor loadings have group-sparse ARD priors

$$[\mathbf{W}_c^{(m)}]_k \sim \mathcal{N}(\mathbf{0}, (\alpha_{c,k}^{(m)})^{-1} \mathbf{I}), \tag{4}$$

where the factors are indexed with $k \in \{1, \ldots, K\}$ and clusters with $c \in \{1, \ldots, S\}$. The shared loadings have similarly $\hat{\mathbf{W}}_k^{(m)} \sim \mathcal{N}\left(\mathbf{0}, \hat{\alpha}_k^{-1} \mathbf{I}\right)$ but here we used different hyperparameters for the ARD. Since the signal is known to be weak, a model that explains all data by shared components would be almost equally good, and therefore we need to tell the model to prefer solutions having some cluster-specific components. That can be achieved by setting the hyperparameter a for the shared components to a larger value, such as $a = 30$ in our experiments. This brings the prior mean to $\langle \alpha \rangle = a/b \gg 1$ while still having a large variance. We recommend setting the value large enough such that during inference both shared and cluster-specific factors are found.

The cluster assignments c_n for each sample n are given by $c_n | \boldsymbol{\pi} \sim \text{Multinomial}(\boldsymbol{\pi})$ with a conjugate Dirichlet prior for the prior probabilities $\boldsymbol{\pi}$. Noise precisions $\tau_c^{(m)}$ have non-informative Gamma priors. The values of τ determine how much of the uninteresting signal can be explained away by the "simple" noise model and how much needs to be explained by the shared components.

2.3 Classifying GFA mixture

We extend the GFA mixture model to include a part that ties the mixture components also to the class labels. Here we consider only two classes, although it is straightforward to extend our model into a multi-class setting. Let r_n be our observed class label (0 or 1) and c_n be latent cluster assignments. We model the output label r_n as a binary variable with a Bernoulli distribution $r_n|c_n, \boldsymbol{\gamma} \sim \text{Bernoulli}(\gamma_{c_n})$, where we use the conjugate prior $\gamma_c \sim \text{Beta}(\frac{1}{2}, \frac{1}{2})$ that denotes the probability $P(r_n = 1|c_n = c)$.

The full joint density for the model is given by

$$p(\Theta|\mathbf{X}, \mathbf{r}) \propto p(\mathbf{X}|\mathbf{W}, \mathbf{Z}, \mathbf{c}, \boldsymbol{\tau})p(\mathbf{r}|\boldsymbol{\gamma}, \mathbf{c})^{\beta}p(\mathbf{c}|\boldsymbol{\pi})p(\boldsymbol{\pi})p(\boldsymbol{\gamma})p(\mathbf{Z})p(\boldsymbol{\tau})p(\mathbf{W}|\boldsymbol{\alpha})p(\boldsymbol{\alpha}), \tag{5}$$

where we have given an additional weight β to the Bernoulli likelihood of the output labels to facilitate better classification results by making the clusters more likely to match the labels. In our experiments we set $\beta = 100$. We recommend to tune the parameter large enough such that the clustering given by the model, starts to match well with the class labels in a training set. Here we also used short-hand notations $\mathbf{X}$ for the data matrix, $\mathbf{W}$ for all loading matrices and $\mathbf{Z}$ for all latent variables, $\mathbf{c}$ for the cluster assignments c_n, and $\mathbf{r}$ for the output labels.

Given the model specification as above, discrimination between the classes is influenced both by the cluster-specific loading matrices $\mathbf{W}_c$ and the noise precisions $\boldsymbol{\tau}_c$ that are allowed to be different for each cluster. The former parameters provide for an easier interpretation by directly inspecting the feature loadings.

2.4 Inference

For inference in our Bayesian model we adopt the Variational Bayesian approach [1], which is based on maximizing a lower bound on the log marginal likelihood of the data for a distribution that is of an easier form than the intractable true posterior distribution. Typically, a factorized approximation $q(\Theta) = \prod_i q(\theta_i)$ is used, where θ_i are some disjoint subsets of variables. It can be shown that the optimal solution then is $q(\theta_i) \propto \exp(\langle \log p(\Theta, \mathbf{X}, \mathbf{r}) \rangle_{\Theta \setminus \theta_i})$ in which the expectation is taken with respect to all variables except θ_i. For our model, we make the factorized approximation $p(\Theta|\mathbf{X}, \mathbf{r}) \approx q(\Theta) = q(\mathbf{c})q(\boldsymbol{\pi})q(\mathbf{Z})q(\hat{\mathbf{W}})q(\hat{\boldsymbol{\alpha}}) \prod_{c=1}^{S} q(\mathbf{W}_c^{(m)})q(\boldsymbol{\tau}_c)q(\gamma^c)q(\boldsymbol{\alpha}_c)$.

3 Experiments with MEG data

MEG data was collected simultaneously from two connected sites, each having a subject in an MEG device and a stimulus presentation computer, a system similar to those by [2] and [23] using Elekta Neuromag 306 channel devices. The pairs of subjects engaged in a word game in which the two subjects took turns in uttering words to come up with a meaningful story. Data are available from 7 pairs of subjects. The lengths of the stories varied between 88–170 words. For our purposes we chose to use data from only one subject of each pair. The data were preprocessed using SSS [17], after which we discarded the magnetometer data leaving only the two gradiometers per sensor location, total of $M = 204$ channels. The data were downsampled from 1000 Hz by a factor of 15 and high-pass filtered at 3 Hz to remove very slow signal changes and drifts.

Fig. 1(a) depicts the conventional ERP analysis, comparing the listening and speaking conditions. We see that particularly in the auditory cortices (located on left and right sides slightly above the center) there exist differences between the two conditions. Differences are also present at central or posterior regions of the scalp.

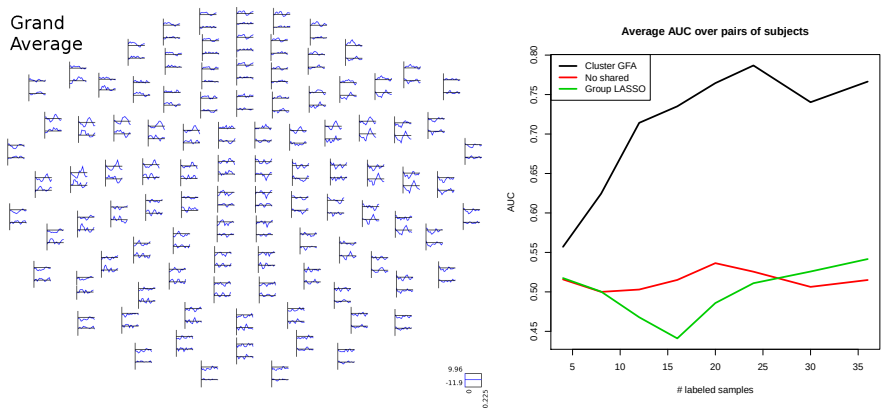

(a) (b)

Fig. 1. (a) Average difference between the "speak" and "listen" conditions over the whole data (all subjects). Each pair of time series describes data of two MEG gradiometers at same location. Nose direction in the figure is upwards. **(b)** Average AUC over the datasets plotted as a function of the size of the training data, for our model with and without the shared factors and group LASSO. The LASSO begins slowly to gain performance only with the biggest sample sizes, while our model copes even with a very small data size.

3.1 Results and interpretation

We are classifying single trials, each trial consisting of either the subject listen a word or speak a word. We consider each MEG data channel as a data source or view. For a comparison method we use a Group LASSO implementation from [3] that sets hyperparameters by cross-validation.

As we have the whole single trial of one MEG gradiometer channel as one data source, the time series structure is modelled through the factor model. More explicitly, the data matrix corresponding to source m is $\mathbf{X}^{(m)} = \left(\mathbf{x}_1^{(m)} \dots \mathbf{x}_N^{(m)}\right)^T$, where N is the number of trials, i.e. the number of spoken words. Each $\mathbf{x}_n^{(m)}$ is a vector of time points within a 300 ms window.

We assess the performance of the classifiers with the AUC statistic and report results based on a resampling technique, where training sets of a given size (see Fig. 1(b)) and test sets of 10 samples were randomly drawn from the full datasets. We drew 10 pairs of training and test sets independently for all subjects; pooling data from multiple subjects is not typically feasible with brain imaging data unless we can first align the data across subjects, see e.g. [4]. In the following results we vary the size of training data N from 4 to 42.

The results are presented in Fig. 1(b) for both our model and group LASSO. Our model is clearly up to the task even with very small training data, whereas group LASSO would require more data to improve performance. Remarkably, a restricted version of our model without shared factors showed very low performance meaning that shared factors were significant. For closer interpretation of the results, we computed our model also for the full datasets. The grand average results in Fig. 1(a) give an idea about which areas of the brain are the most discriminative. As our model is generative, we can generate cluster-specific ERPs from the estimated model. The reconstruction of channel m is calculated for cluster c as $\hat{\mathbf{X}}^{(m)} = \mathbf{Z}\mathbf{W}_c^{(m)T}$ from which we average over the trials to obtain the ERPs shown in Fig. 2. We found that the reconstructions and the averages computed directly from

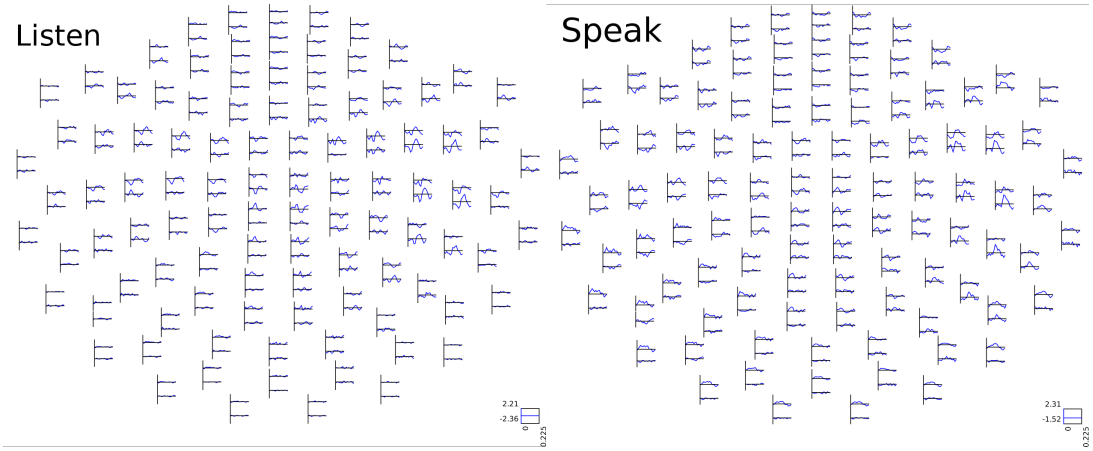

Fig. 2. Reconstructed data from factors in classes "listen" and "speak," and the shared factors. The difference shows what parts of the total response are involved in explaining the class differences.

data matched closely; the model found the existing differences and correctly picked them into the class-related clusters. With smaller training sets, the reconstructions were partial as fewer number of factors were in the model.

This simple case study demonstrated that the model is able to find discriminative signals even in single-trial MEG data. In this specific data the most discriminative signals are likely related to muscle activity which is present during speaking but not when listening; thus the model has picked up signs of this activity which is typically most visible in the channels closest to the edges of the MEG helmet. Also other areas are active in the "speak" condition. For the "listen" condition, both our model and the grand average show activations clearly in horizontally central areas that include the auditory cortex on both sides responsible for processing the word spoken by the other participant.

4 Conclusion

We introduced a classification model based on a mixture of group factor analyzers that share some of their factors. The sharing seemed to be very significant in improving classification accuracy on our brain imaging datasets; the model without shared factors performed much worse, as did the group LASSO baseline. In addition, we showed that our model gives interpretable results. The proposed model included two parameters that required tuning, namely the hyperprior parameter for the precisions of the shared loading matrices, and the weighting parameter for the Bernoulli distribution of the class labels. We provided a simple way to set these, but it would be preferable to have a more rigorous analysis in future work.

Acknowledgments. This work was financially supported by MindSEE (FP7 – ICT; Grant Agreement #611570) and the Academy of Finland (CoE in Computational Inference Research COIN and LASTU). The dual-MEG data set was collected in Brain2Brain project funded by the European Research Council (Advanced Grant #232946 to Riitta Hari, Brain Research Unit, O.V. Lounasmaa Laboratory, Aalto University). We thank P. Baess, R. Hari, T. Himberg, L. Hirvenkari, V. Jousmäki, A. Mandel, J. Mäkelä, J. Nurminen, L. Parkkonen, and A. Zhdanov for the possibility to use the anonymized

dual-MEG data, A. Mandel for preprocessing of the data set, and L. Hirvenkari for the stimulus timing. Computational resources were provided by Aalto Science-IT project.

References

1. Attias, H.: Inferring parameters and structure of latent variable models by variational Bayes. In: Proceedings of the 15th Conference on Uncertainty in Artificial Intelligence. pp. 21–30. Morgan Kaufmann Publishers Inc. (1999)
2. Baess, P., Zhdanov, A., Mandel, A., Parkkonen, L., Hirvenkari, L., Mäkelä, J.P., Jousmäki, V., Hari, R.: MEG dual scanning: a procedure to study real-time auditory interaction between two persons. Frontiers in Human Neuroscience 6 (2012)
3. Breheny, P., Huang, J.: Group descent algorithms for nonconvex penalized linear and logistic regression models with grouped predictors. Statistics and Computing 25(2), 173–187 (2015)
4. Conroy, B.R., Singer, B.D., Guntupalli, J.S., Ramadge, P.J., Haxby, J.V.: Inter-subject alignment of human cortical anatomy using functional connectivity. NeuroImage 81, 400–411 (2013)
5. Ek, C.H., Torr, P.H., Lawrence, N.D.: Gaussian process latent variable models for human pose estimation. In: Machine Learning for Multimodal Interaction, pp. 132–143. Springer (2008)
6. Eleftheriadis, S., Rudovic, O., Pantic, M.: Shared Gaussian process latent variable model for multi-view facial expression recognition. In: Advances in Visual Computing, pp. 527–538. Springer (2013)
7. Haufe, S., Meinecke, F., Görgen, K., Dähne, S., Haynes, J.D., Blankertz, B., Biessmann, F.: Parameter interpretation, regularization and source localization in multivariate linear models. In: Proceedings of the 4th International Workshop on Pattern Recognition in Neuroimaging (PRNI). pp. 1–4. IEEE (2014)
8. Höhne, J., Blankertz, B., Muller, K.R., Bartz, D.: Mean shrinkage improves the classification of ERP signals by exploiting additional label information. In: Proceedings of the 4th International Workshop on Pattern Recognition in Neuroimaging (PRNI). pp. 1–4. IEEE (2014)
9. Jia, Y., Salzmann, M., Darrell, T.: Factorized latent spaces with structured sparsity. In: Advances in Neural Information Processing Systems, vol. 23, pp. 982–990 (2010)
10. Kan, M., Shan, S., Zhang, H., Lao, S., Chen, X.: Multi-view discriminant analysis. In: Computer Vision–ECCV 2012, pp. 808–821. Springer (2012)
11. Kia, S.M., Vega-Pons, S., Olivetti, E., Avesani, P.: Multi-task learning for interpretation of brain decoding models. In: MLINI 2014 - 4th NIPS Workshop on Machine Learning and Interpretation in Neuroimaging (2014)
12. Klami, A., Virtanen, S., Kaski, S.: Bayesian canonical correlation analysis. Journal of Machine Learning Research 14, 965–1003 (2013)
13. Klami, A., Virtanen, S., Leppäaho, E., Kaski, S.: Group factor analysis. IEEE Transactions on Neural Networks and Learning Systems pp. 26(9):2136–2147 (2015)
14. Lemm, S., Blankertz, B., Dickhaus, T., Müller, K.R.: Introduction to machine learning for brain imaging. NeuroImage 56(2), 387–399 (2011)
15. Reynolds, D.A.: Speaker identification and verification using Gaussian mixture speaker models. Speech Communication 17(1), 91–108 (1995)
16. Santana, R., Mendiburu, A., Lozano, J.A.: Multi-view classification of psychiatric conditions based on saccades. Applied Soft Computing 31, 308–316 (2015)
17. Taulu, S., Kajola, M., Simola, J.: Suppression of interference and artifacts by the signal space separation method. Brain Topography 16, 269–275 (2004)
18. Virtanen, S., Klami, A., Khan, S.A., Kaski, S.: Bayesian group factor analysis. In: Lawrence, N., Girolami, M. (eds.) Proceedings of the 15th International Conference on Artificial Intelligence and Statistics. JMLR W&CP, vol. 22, pp. 1269–1277. JMLR (2012)
19. Wang, H., Nie, F., Huang, H.: Multi-view clustering and feature learning via structured sparsity. In: Proceedings of the 30th International Conference on Machine Learning (ICML). pp. 352–360 (2013)

20. Yuan, M., Lin, Y.: Model selection and estimation in regression with grouped variables. Journal of the Royal Statistical Society: Series B (Statistical Methodology) 68(1), 49–67 (2006)
21. Zhang, H., Nasrabadi, N.M., Zhang, Y., Huang, T.S.: Joint dynamic sparse representation for multi-view face recognition. Pattern Recognition 45(4), 1290–1298 (2012)
22. Zhao, S., Gao, C., Mukherjee, S., Engelhardt, B.E.: Bayesian group latent factor analysis with structured sparse priors. JMLR, to appear (2015), pre-print arXiv:1411.2698
23. Zhdanov, A., Nurminen, J., Baess, P., Hirvenkari, L., Jousmäki, V., Mäkelä, J.P., Mandel, A., Meronen, L., Hari, R., Parkkonen, L.: An internet-based real-time audiovisual link for dual MEG recordings. PLOS ONE 10(6), e0128485 (2015)